

Using Genetic Algorithms for Weights Space Exploration in an Eurovision-like weighted Meta-Classifier

Radu G. Crețulescu, Daniel I. Morariu, Macarie Breazu and Lucian N. Vințan

Abstract — Automatic document classification has become an important task because of the continually increasing number of text documents with the users have to deal. The aim of this paper is to develop a non-adaptive meta-classifier for text documents in order to increase the classification accuracy. The developed meta-classifiers are based on combining some SVM classifiers and a Naïve Bayes classifier. We proposed a new meta-classification method which takes into consideration the corresponding positions and confidence degrees obtained for all the classes. In this work we have tried to find, using Genetic Algorithms, the optimal weighting factors for the values returned by each classifier separately. Consequently, there is a possibility for the meta-classifier to decide that the winner is a class that is not necessarily one of the winner classes returned by the individual classifiers. The experimental results have showed that the classification accuracy can be improved through the proposed method.

Keywords— Text Classification and Performance Evaluation, SVM, Meta-classification, Genetic Algorithms

1 INTRODUCTION

WHILE more and more textual information is available online, effective retrieval is difficult without good indexing and summarization of document content. Document categorization is one solution to this problem. The task of document categorization is to assign a user defined categorical label to a given document. In recent years a growing number of categorization methods and machine learning techniques have been developed and applied in different contexts.

Documents are typically represented as vectors in a features space. Each word in the vocabulary is represented as a separate dimension. The number of a certain word's occurrences in a document represents the value of the corresponding component in the document's vector.

In this paper we investigate a strategy for combining classifiers' results in order to improve the classification accuracy using genetic algorithms. We used classifiers based on Support Vector Machine (SVM) techniques and based on Naïve Bayes theory. They are less vulnerable to degrade with an increasing dimensionality of the feature space, and have been shown effective in many classification tasks. The SVM classifier is actually based on learning with kernels and support vectors.

We combine multiple classifiers hoping that the classification accuracy can be improved without a significant increase in response time. Instead of building just one highly accurate specialized classifier with much time and effort, we build and combine several simpler classifiers.

Several combination schemes have been described in the literature [4] and [7]. A usual approach is to build individual classifiers and later combine their judgments to make the final decision. Another approach, which is not so commonly used because it suffers from the “curse of dimensionality” [6], is to concatenate features from each classifier to make a longer feature vector and use it for the final decision. Anyway, meta-classification is effective only if its component classifiers synergies can be exploited.

In previous studies combination strategies were used ad hoc and are strategies like majority vote, linear combination, winner-take-all [4], or Bagging and Adaboost [15]. Also, some rather complex strategies have been suggested. For example in [6] and [9] meta-classification strategies using SVM [14] are presented and are compared with probability based strategies.

Section 2 and 3 contains prerequisites for the main work developed in this research. In sections 4 we present the methodology used for developing our experiments. Section 5 presents the experimental framework and section 6 presents the main results of our experiments. Finally the last section debates and concludes on the most important results obtained and proposes some further work.

2 CLASSIFIERS USED

2.1 Support Vector Machine

The Support Vector Machine (SVM) is a classification technique based on statistical learning theory that was applied with great success in many challenging non-linear classification problems and on large data sets ([12], [14]).

The SVM algorithm finds a hyperplane that optimally splits the training set. The optimal hyperplane can be distinguished by the maximum margin of separation between all training points and itself. Looking at a two-dimensional problem we actually want to find a line that “best” separates points in the positive class from points in the negative class. The hyperplane is characterized by a decision function like:

$$f(x) = \text{sgn}(\langle \mathbf{w}, \Phi(x) \rangle + b) \quad (1)$$

where \mathbf{w} is the weight vector, orthogonal to the hyperplane, “ b ” is a scalar that represents the margin of the hyperplane, “ x ” is the current sample tested, “ $\Phi(x)$ ” is a function that transposes the input data into a higher dimensional feature space and $\langle \cdot, \cdot \rangle$ representing the dot product. Sgn is the sign function. If \mathbf{w} has unit length, then $\langle \mathbf{w}, \Phi(x) \rangle$ is the length of $\Phi(x)$ along the direction of \mathbf{w} . Generally \mathbf{w} will be scaled by $\|\mathbf{w}\|$. In the training part the algorithm needs to find the normal vector “ \mathbf{w} ” that leads to the largest “ b ” of the hyperplane.

2.2 Naïve Bayes

The Bayes classifier use the Bayes Theorem which basically computes prior probabilities for a given class based on the probability for a given term to belong to the specified class. Thus the classifier computes the probability for a document to be into a given class.

Bayesian theory works as a framework for making decision under uncertainty - a probabilistic approach to inference [5] - and it is particularly suited when the dimensionality of the inputs data is

high. Bayes theorized that the probability of future events could be calculated by determining their earlier frequency.

The Naive Bayes classifier is based on the simplified assumption that the attribute values are conditionally independent given target value. In other words the assumption is that, given the target value of the instance, the probability of observing the conjunction y_1, y_2, \dots, y_n is just the product of the probabilities for the individual attributes:

$$c_{map} = \arg \max_{1 < i < m} \overline{P}(X_i|Y) = \arg \max_{1 < i < m} \overline{P}(X_i) \prod_{j=1}^n \overline{P}(y_j|X_i), \quad (2)$$

For extending the SVM and the Naïve Bayes classifiers from two-class classification to multi-class classification typically one of two methods are used: “One versus the rest”, where each topic is separated from the remaining topics, and “One versus the one”, where a separate classifier is trained for each class pair [14]. We selected the first method for two reasons. First, preliminary experiments show that this method gives better performance, which might be explained by the fact that the Reuter’s database contains strongly overlapped classes and assigns almost all samples in more than one class. Second, overall training time is much shorter for the first method.

3 GENETIC ALGORITHMS

Genetic algorithms encode a potential solution to a specific problem on a simple chromosome-like data structure and apply genetic operators to these structures in order to preserve critical information [16]. In our weights selection problem the chromosome is considered to be of the following form:

$$c = (w_1, w_2, \dots, w_k) \quad (3)$$

where $w_i, i = \overline{1, k}$ represent the weight for each class position returned by every classifier used into the meta-classifier. The training set has the form $\{\langle \bar{x}_{ij}, y \rangle, i = \overline{1, m} \text{ and } j = \overline{1, n}\}$, where y represents the correct output for the input sample \bar{x}_{ij} , n represents the number of classes and m represents the number of classifiers used into the meta-classifier. For the genetic algorithm the training dataset consist of all outputs from all classifiers used into the meta-classifier. For simplifying the computing of the fitness function we choose a representation of the chromosome that uses the outputs of all classifiers from the meta-classifier. With this approach we take into consideration the position of each class in the classifier’s output. The chromosome keeps only, for each class position, the weights that are used to compute the accuracy of meta-classifier. Thus, the potential solution of the problem encodes the weights that can be used to compute the final decision into meta-classifier [11].

For each chromosome we can compute the fitness function as the final classification accuracy of the meta-classifier using the testing file. The classification accuracy is computed as the number of final correct classified classes by the meta-classifier.

For each test we start with a generation of 100 chromosomes, each of them having values randomly generated between -1 and 1. In the next step we generate the next population using genetic operators as: selection, crossover or mutation [16]. The evolutionary process stops after a

predefined number of 1000 generations. The evolutionary process stops after a predefined number of generations are taken or when in the last 20 generations no changes occur.

At the end of the algorithm, we obtain for each class position the “best” weight that can be used for computing the winner class. The general scheme of the genetic algorithm is presented in pseudo code in the Algorithm 1:

```
Begin
For each topic from a topics_set
  begin
    generate a population
    while not terminated condition
      For each chromosome from population
        compute the fitness functions
        make next population:
          select parents
          recombine pairs of parents
          apply mutation to offspring
      End while.
      Store the chromosome that obtain the best fitness
    End for.
    Take all stored chromosomes
End.
```

Algorithm 1: Pseudo code for GA algorithm

4 META-CLASSIFIER MODELS

In our previous work [2], [8] and [10] it is presented a meta-classifier, based on 8 SVM classifiers and one Bayes classifier, that were used to improve the classification accuracy for text documents. In those works 3 meta-classifier models are used: majority vote, selection based on the Euclidian distance and selection based on the cosine angle. The first model was a non-adaptive model that had the advantage of speed, but obtained not so good results. The last two were adaptive models that used a training part before they can be used in the classification. The training part is much time consuming and, unfortunately, this time increases when the meta-classifier learn more examples.

In this paper we propose a new non-adaptive meta-classifier model based on some pre-optimized weights that will increase the classification accuracy. In [11] we have presented a solution for this problem that uses static (pre-determined) values for the weights. Now we propose to use a genetic algorithm for pre-computing the optimal weights. Obviously, this adaptive optimization process will be done before the meta-classification process will start. The obtained weights' values will correspondingly weight all the results returned by every classifier for each document, as it is shown further.

All meta-classifiers presented in this article contain eight SVM type classifiers and one naïve Bayes classifier.




Fig. 4.1 Creating entry data for meta-classifiers

4.1 Non-adaptive Meta-classifier based on Majority Vote (M-MV)

This meta-classifier model, also presented in [8], is a non-adaptive model that obtains the same results for the same input every time. The idea is to use all selected classifiers to classify the current document. Each classifier will propose a class for the given document, and the meta-classifier will increment the corresponding class-counter. The winner class proposed by the meta-classifier is the class with the highest count. If we obtain two or more classes with identical maximal counts then we classify the current document in all the proposed classes. This method uses only the winner class from the classification returned by each classifier separately. The classes found on the subsequent positions are not taken into consideration.

4.2 Non-adaptive Meta-classifier based on sum (M-SUM)

Each classifier has as input a document that is represented as a vector having 1306 features and produces an output vector with 16 values. The output values represent the confidence degrees, given by the classifier. These values represent the “belonging degree” of the current document to each one of the 16 classes (see Fig. 4.1). More precisely, each element from the output vector represents the classifier’s decision function value for each class separately. So far, in [8], as it was already described in paragraph 4.1, the highest value from the output vectors has been chosen and the class corresponding to that value was considered to be the winner class. The methods proposed in the following paragraph, take into account both the value obtained for each class separately and the corresponding rating position. Since there are 9 classifiers in the meta-classifier for each document there will be 9 corresponding vectors, each of them having 16 values (so-called confidence degrees).

The values of the decision functions for the SVM type classifiers are in the range $(-\infty, \infty)$ but usually close to the value 0, and the values of the Naive Bayes classifier are in the range $(-\infty, 0)$. Considering these differences, and for making a correct sum of those values, we have transposed the values of all vectors into the interval $[1, \infty)$. The formula used for transposing the values is:

$$V'_i = V_i + \left| \min(\vec{V}) \right| + 1, \text{ for } i = \overline{1,16} \quad (4)$$

This transposing formula doesn’t change the differences between the elements values from a vector even if it contains positive and negative values or only positive ones. Thus, for each vector the differences between its values remain unchanged. In order to calculate the sum of these vectors in the next step we have normalized the vectors bringing their values in the $(0, 1]$ interval. This

normalization ensures that the value from the first position, of the descended arranged vector, is always 1. We avoid also that the value from the last position to be 0.

$$V'_i = \frac{V_i}{\max(\vec{V})} \quad (5)$$

In the current meta-classifier, that only makes the un-weighted sums, (called here M-SUM), we have calculated the sum of corresponding scalars from these 9 vectors. Thus we have obtained, for the current document, a single vector with 16 values. The meta-classifier will decide the winning class as the class with the highest obtained value. If there are two or more classes with identical maximal values, the meta-classifier will propose all those classes.

$$Class = \max_{c, i=1,16} \sum_{k=1}^9 V_i[k] \quad (6)$$

This approach gives the possibility that, even if a class that was never selected as a winner by the classifiers but always obtained a value very close to the winner class, such a class should have the chance that, after summing, to be selected by the meta-classifier as the winner class. The idea started from the fact that in many articles like [3], [2] and [10] we observed the tendency to use methods for choosing the second or the third class as the winner class mostly when it was obvious that the meta-classifier would not provide the correct winner class.

4.3 Non-adaptive Meta-classifier based on weighted sum (M-GSUM)

In this section we will introduce a new non-adaptive meta-classifier based on the weighted sum (we will call it as M-WSUM).

In this meta-classifier the values for the classes will be multiplied with a value, called weight. In [11] are presented some meta-classifiers using different weight values and their obtained results. Since choosing the optimum weight values for each class position is a difficult task, in this article we propose a method based on a genetic algorithm for computing the weights.

Starting from the test dataset [11] we have created a proper test file which is used for the genetic algorithm. In this new file, for each document entry we have saved all outputs from all classifiers separately. Thus for a document we obtained as output 9 vectors (because we have 9 classifiers), each vector having a number of 16 scalars (because we have 16 classes). In the genetic algorithm used one chromosome represents the weights value for each class separately, depending by the class position. The first value from the vector represents the weight of the first position class from the classifier; the second value from the vector represents the weight of the class from the next position, and so on.

$$c = (w_1, w_2, \dots, w_j), \text{ with } j = \overline{1,16} \quad (7)$$

where w_j represents the value of weight for the j class position. We have fulfilled the following inequalities: $w_1 > w_2 > \dots > w_{16}$. For computing the fitness function the following steps are taken:

1. For each document from the testing set the algorithm obtains the output vector corresponding to each classifier;
2. Each output vector is weighted with the corresponding chromosome value;

3. The winner class is obtained by summing all results for each class and selecting the highest score (formula 8);

$$Class_j = \sum_{i=1}^9 w_j c_{ij}, \text{ for } j = \overline{1,16} \quad \text{and} \quad WinClass = \arg \max_{j=1,16}(Class_j) \quad (8)$$

where w_j represents the value of the chromosome for the j^{th} class position (called weight), c_{ij} represents the value computed by the i^{th} classifier for the j^{th} class position (called confidence degree) and $Class_j$ represents the value computed for the class from the j^{th} position.

4. The winner class ($WinClass$) will be compared with the real class proposed by Reuters and if the classes are identically we consider that the document was correctly classified;
5. After processing all documents the fitness value is the accuracy obtained on the whole testing set.

The best chromosome is that with the highest fitness. For building a population we have used 100 chromosomes and the algorithm was applied for 1000 generations.

As selection operator for choosing the chromosome from the current population to be used for the next population we have used two methods: Roulette method and Gaussian method. For the Gaussian method we have used the following formula:

$$Gauss(c) = e^{-\frac{1}{2} \frac{(fitness(c)-m)^2}{\sigma^2}} \quad (9)$$

where c is the current chromosome, $fitness(c)$ is the result obtained for the corresponding chromosome c , m represents the average and in our case is equal to 1 and σ is the standard deviation (here we have used a value equal with 0.5).

For creating the new population we have used all 3 genetic operators thus: 30% for new population is created using the selection operator (selecting the best chromosome - elitism), the rest of the population is created by using the mutation operator in 40% and crossover operator in 30%. After applying this operators the condition $w_1 > w_2 > \dots > w_{16}$ is respected. Thus, for crossover operator we search after a cut point (i) for which the condition $w_i^{first_parent} > w_{i+1}^{second_parent}$ is valid. For mutation operator, after selecting randomly the mutation point (i) the new value is selected randomly in the interval $w_i \in (w_{i-1}, w_{i+1})$.

5 EXPERIMENTAL FRAMEWORKS

The Dataset

Our experiments were performed on the Reuters-2000 collection [13], which has 984 Mb of newspapers articles in a compressed format. The collection includes a total of 806,791 documents, with news stories published by Reuters Press covering the period from 20.07.1996 through 19.07.1997. The articles have 9822391 paragraphs and contain 11522874 sentences and 310033 distinct root words. Documents are pre-classified according to 3 categories: by the Region (366 regions) the article refers to, by Industry Codes (870 industry codes) and by Topics proposed by Reuters (126 topics, 23 of them contain no articles). Due to the huge dimensionality of the database we will present here results obtained using a subset of data. From all documents we selected the

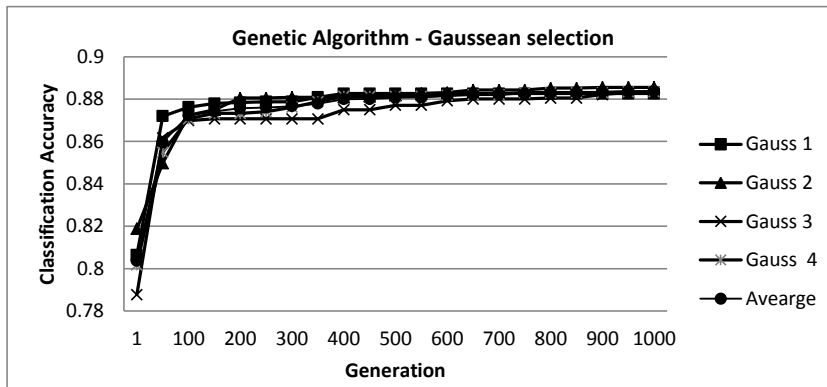


Fig.6.1 Classification accuracy - Gaussian method for selecting the best chromosomes for new generation

documents for which the industry code value is equal to “System software”. We obtained 7083 files that are represented using 19038 features and 68 topics. We have represented a document as a vector of words, after applying a stop-word filter (from a standard set of 510 stop-words) and extracting the word stem [1]. From these 68 topics we have eliminated those topics that are poorly or excessively represented. Thus we eliminated those topics that contain less than 1% documents from all 7083 documents in the entire set. We have also eliminated topics that contain more than 99% samples from the entire set, as being excessively represented. After doing so we have obtained 24 different topics and 7053 documents, that were split randomly in a training set (4702 samples) and a testing set (2351 samples). In the feature extraction phase we take into consideration both the article and the title of the article. In the feature selection phase we have selected only 1306 features for each vector.

6 EXPERIMENTAL RESULTS

In [11] were presented first time the results obtained using static (pre-determined) weights values. The best classification accuracy obtained was 87.20% (i.e. 301 incorrectly classified documents

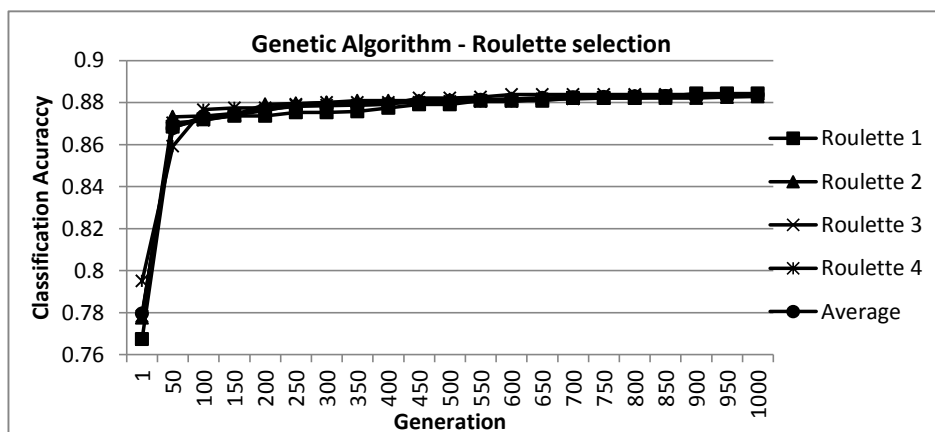


Fig. 6.2 Classification accuracy with Roulette method for selecting the best chromosome for new generation

from a number of 2351 documents). These results were obtained with a model of meta-classifier called “M-05W” that has weighting values that linearly decrease from 12 with a step of 0.5.

In the meta-classifier with genetic algorithm, for each chromosome we compute the fitness function using the training dataset. After selecting the best chromosome in this manner, we compute the fitness function obtained only by this chromosome using the testing dataset and in the figures only the values computed on the testing dataset are presented. Therefore, sometimes, the evolution of the best chromosome in the presented figures is not always ascending.

Because we start with a randomly generated population, and the obtained results depend of the starting points, we make four successively runs for the same method. Here we present the output for each method and also the average value.

In figure 6.1 we present the classification accuracy obtained on the testing dataset, when using the genetic algorithm with Gaussian method for selecting the chromosomes. The average accuracy obtained is 88.37%. Thus we obtained an average growth of 1.17% for the classification accuracy comparing with the case of the static values of weights presented in [11].

In case of using genetic algorithm with the Roulette method for selecting the chromosomes in the new population we have obtained an average classification accuracy of 88.36%, with only 1.16% greater the in M-SUM method. The results are presented in figure 6.2.

The best accuracy obtained by us in all four performed tests was obtained using Gaussian method for chromosome selection where the maximum accuracy value was 88.55%.

7 CONCLUSIONS AND FURTHER WORK

In this article we presented a non-adaptive meta-classifier used for classifying text documents. This meta-classifier uses the outputs of eight independent SVM classifiers and a Naive Bayes classifier. The output of each classifier is a vector of values, where each position represents the confidence given by the classifier that the current document belongs to the corresponding class. The meta-classifier developed in this article uses a genetic algorithm for calculating the best values that could be used to weight the outputs of each classifier so that the final classification accuracy to be improved.

Since we started for the genetic algorithm with random initial weights we conducted four separate runs. As the average results compared with results obtained when classifiers outputs were weighted with pre-defined values [11] we have obtained an improvement of 1.17%. As classification accuracy, the best value obtained is 88.55% in case of using Gaussian method for selection of chromosomes for the new population.

In the further, an interesting natural extension of our work may be an adaptive and intelligent meta-classifier that uses a neural network for choosing the classifier that will be used in classifying the current document.

8 ACKNOWLEDGMENT

This work was partially supported by CNCSIS-UEFISCSU, project number PN II-RU code PD_670/2010.

9 REFERENCES

- [1] S. Chakrabarti, *Mining the Web- Discovering Knowledge from hypertext data*, Morgan Kaufmann Press, 2003.
- [2] R. Cretulescu, *Support Vector Machine versus Bayes Naïve*, 2nd PhD Report, “Lucian Blaga” University of Sibiu, 2008, <http://webspaces.ulbsibiu.ro/radu.kretzulescu/html/phdreport2.pdf>
- [3] Q. Chen, D. Zheng, T. Zhao, S. Li, *A Fusion of Multiple Classifiers Approach Based on Reliability Function for Text Categorization*, 5th International Conference on Fuzzy Systems and Knowledge Discovery, IEEE, 2008.
- [4] N. Dimitrova, L. Agnihotri and G. Wei, *Video Classification Based on HMM Using Text and Face*, Proceedings of the European Conference on Signal Processing, Finland, 2000.
- [5] D. Lewis, *Naive (Bayes) at Forty: The Independence Assumption in Information Retrieval*. In Proceedings of the 10th European Conference on Machine Learning, 1998
- [6] W.-H. Lin, A. Hauptmann, *News Video Classification Using SVM-based Multimodal Classifier and Combination Strategies*, In Proceedings of the Tenth ACM international Conference on Multimedia, 2002
- [7] W.-H. Lin, R. Jin, A. Hauptmann, *A Meta-classification of Multimedia Classifiers*, International Workshop on Knowledge Discovery in Multimedia and Complex Data, Taiwan, 2002.
- [8] D. Morariu, L. Vintan, V. Tresp, *Meta-Classification using SVM Classifiers for Text Documents*, The 3rd International Conference on Neural Computing and Pattern Recognition, Barcelona, October 2006.
- [9] D. Morariu, *Text Mining Methods based on Support Vector Machine*, MatrixRom, Bucharest, 2008.
- [10] D. Morariu, R. Cretulescu, L. Vintan, *Improving a SVM Meta-Classifier for Text Documents by using Naïve Bayes*, International Journal of Computers, Communications & Control, ISSN 1841-9836, E-ISSN 1841-9844, 2010.
- [11] R. Cretulescu D. Morariu, L. Vintan, *Eurovision-like weighted Non-Adaptive Meta-classifier for Text Documents*, The 8th RoEduNet International Conference, Galati, Romania, 2009.
- [12] C. Nello, J. Swawe-Taylor, *An introduction to Support Vector Machines*, Cambridge University Press, 2000.
- [13] Reuters Corpus: <http://about.reuters.com/researchandstandards/corpus/>. Released in November 2000.
- [14] B. Schoelkopf, A. Smola, *Learning with Kernels. Support Vector Machines*, MIT Press, London, 2002.
- [15] G. Siyang, L. Quingrui, M. Lin, *Meta-classifier in Text Classification*, <http://www.comp.nus.edu.sg/~zhouyong/papers/cs5228project>.
- [16] A. Ghosh, L.C. Jain, *Evolutionary Computation in Data Mining*, Springer Verlag Berlin Heidelberg, 2005